Refine Search

Search Results -

Terms	Documents
L11 and @pd > 20060816	1

Database: US Pre-Grant Publication Full-Text Database; US Patents Full-Text Database; US OCR Full-Text Database; EPO Abstracts Database; JPO Abstracts Database; Derwent World Patents Index; IBM Technical Disclosure Bulletins

Search: L12

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, December 05, 2006 **Purge Queries** Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ		
L12	L11 and @pd > 20060816	1	L12
L11	((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and (source or target adj table)and (captur$ same information))	10	L11
L10	((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and (source or target adj table)and (capt$ same information))	10	L10
L9	711/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and (source or target adj table))	0	L9
L8	709/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and (source or target adj table))	1	L8
	707/$.ccls. and ((replicat$ same data same (change$ or updat$) same		

L7	asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and ((source or target adj table) with (key adj column)) and stor$)	2	L7
L6	707/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and ((source or target adj table) with (key adj column)) and stro$)	0	L6
L5	707/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and ((source or target adj table) with (key adj column)))	2	L5
L4	707/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$) and (source or target adj table))	14	L4
L3	707/$.ccls. and ((replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$))	15	L3
L2	(replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log") and (build$ or creat$ or generat$ same quer$ or search$)	17	L2
L1	(replicat$ same data same (change$ or updat$) same asynchron$)and (database adj log or "database log")	18	L1

END OF SEARCH HISTORY

Terms used
asynchronously replicating data database log data value source table

Found **7,323** of **193,448**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown

Relevance scale

1 Replication: DB2, Oracle, or Sybase?

Doug Stacey
December 1995 **ACM SIGMOD Record**, Volume 24 Issue 4
Publisher: ACM Press
Full text available: pdf(726.69 KB) Additional Information: full citation, abstract, citings, index terms

Is replication salvation or the devil in disguise? Here's what three implementations tell us

2 Distributed systems - programming and management: Optimistic replication in HOPE

Crispin Cowan
November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2**
Publisher: IBM Press
Full text available: pdf(895.21 KB) Additional Information: full citation, abstract, references

The growing imbalance between network latency and throughput is causing nodes in distributed systems to appear to be moving farther apart. Many distributed systems are turning to replication as a mechanism to make resources appear closer. However, maintaining one-copy consistency in a system containing replicated elements has proven difficult. In particular, the checks required to make consistent updates can take as long or longer than the time that would be taken to update a single remote copy. ...

3 Performance evaluation of extended storage architectures for transaction processing

Erhard Rahm
June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2
Publisher: ACM Press
Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings, index terms

The use of non-volatile semiconductor memory within an extended storage hierarchy promises significant performance improvements for transaction processing. Although page-addressable semiconductor memories like extended memory, solid-state disks and disk caches are commercially available since several years, no detailed investigation of their use for transaction processing has been performed so far. We present a comprehensive simulation study that compares the performance of these storage ty ...

4

Garbage collection for a client-server persistent object store
Laurent Amsaleg, Michael J. Franklin, Olivier Gruber
August 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 3
Publisher: ACM Press

Full text available: pdf(267.18 KB) Additional Information: full citation, abstract, references, citings, index terms, review

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environmnet. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does non require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

5 The state of the art in automating usability evaluation of user interfaces
Melody Y. Ivory, Marti A Hearst
December 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 4
Publisher: ACM Press

Full text available: pdf(2.31 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Usability evaluation is an increasingly important part of the user interface design process. However, usability evaluation can be expensive in terms of time and human resources, and automation is therefore a promising way to augment existing approaches. This article presents an extensive survey of usability evaluation methods, organized according to a new taxonomy that emphasizes the role of automation. The survey analyzes existing techniques, identifies which aspects of usability evaluation aut ...

Keywords: Graphical user interfaces, taxonomy, usability evaluation automation, web interfaces

6 Application level performance: On the use and performance of content distribution networks
Balachander Krishnamurthy, Craig Wills, Yin Zhang
November 2001 **Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement**
Publisher: ACM Press

Full text available: pdf(2.51 MB) Additional Information: full citation, abstract, references, citings, index terms

Content distribution networks (CDNs) are a mechanism to deliver content to end users on behalf of origin Web sites. Content distribution offloads work from origin servers by serving some or all of the contents of Web pages. We found an order of magnitude increase in the number and percentage of popular origin sites using CDNs between November 1999 and December 2000.In this paper we discuss how CDNs are commonly used on the Web and define a methodology to study how well they perform. A performanc ...

7 Experience with transactions in QuickSilver
Frank Schmuck, Jim Wylie
September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5

Publisher: ACM Press

Full text available: pdf(1.66 MB) Additional Information: full citation, abstract, references, citings, index terms

All programs in the QuickSilver distributed system behave atomically with respect to their updates to permanent data. Operating system support for *transactions* provides the framework required to support this, as well as a mechanism that unifies reclamation of resources after failures or normal process termination. This paper evaluates the use of transactions for these purposes in a general purpose operating system and presents some of the lessons learned from our experience with a complet ...

8 Classification - 1: Multi-evidence, multi-criteria, lazy associative document classification

Adriano Veloso, Wagner Meira, Marco Cristo, Marcos Gonçalves, Mohammed Zaki
November 2006 **Proceedings of the 15th ACM international conference on Information and knowledge management CIKM '06**
Publisher: ACM Press
Full text available: pdf(269.48 KB) Additional Information: full citation, abstract, references, index terms

We present a novel approach for classifying documents that combines different pieces of evidence (e.g., textual features of documents, links, and citations) transparently, through a data mining technique which generates rules associating these pieces of evidence to predefined classes. These rules can contain any number and mixture of the available evidence and are associated with several quality criteria which can be used in conjunction to choose the "best" rule to be applied at classification t ...

Keywords: classification, data mining, lazy algorithms

9 Data integration and sharing II: Scientific data repositories: designing for a moving target

Etzard Stolte, Christoph von Praun, Gustavo Alonso, Thomas Gross
June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**
Publisher: ACM Press
Full text available: pdf(739.27 KB) Additional Information: full citation, abstract, references, citings, index terms

Managing scientific data warehouses requires constant adaptations to cope with changes in processing algorithms, computing environments, database schemas, and usage patterns. We have faced this challenge in the RHESSI Experimental Data Center (HEDC), a datacenter for the RHESSI NASA spacecraft. In this paper we describe our experience in developing HEDC and discuss in detail the design choices made. To successfully accommodate typical adaptations encountered in scientific data management systems ...

10 Higher-order distributed objects

Henry Cejtin, Suresh Jagannathan, Richard Kelsey
September 1995 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 17 Issue 5
Publisher: ACM Press
Full text available: pdf(2.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are user-specified load-balancing and migration policies for threads, incrementally linked

distributed computations, and parameterized client-server applicat ...

Keywords: concurrency, continuations, higher-order languages, message-passing

11 Network performance effects of HTTP/1.1, CSS1, and PNG

Henrik Frystyk Nielsen, James Gettys, Anselm Baird-Smith, Eric Prud'hommeaux, Håkon Wium Lie, Chris Lilley

October 1997 **ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '97**, Volume 27 Issue 4

Publisher: ACM Press

Full text available: pdf(1.62 MB) Additional Information: full citation, abstract, references, citings, index terms

We describe our investigation of the effect of persistent connections, pipelining and link level document compression on our client and server HTTP implementations. A simple test setup is used to verify HTTP/1.1's design and understand HTTP/1.1 implementation strategies. We present TCP and real time performance data between the libwww robot [27] and both the W3C's Jigsaw [28] and Apache [29] HTTP servers using HTTP/1.0, HTTP/1.1 with persistent connections, HTTP/1.1 with pipelined requests, and ...

12 Parallel database processing on a 100 Node PC cluster: cases for decision support query processing and data mining

Takayuki Tamura, Masato Oguchi, Masaru Kitsuregawa

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)**

Publisher: ACM Press

Full text available: pdf(157.74 KB) Additional Information: full citation, abstract, references, citings

We developed a PC cluster system consists of 100 PCs. Each PC employs the 200MHz Pentium Pro CPU and is connected with others through an ATM switch. We picked up two kinds of data intensive applications. One is decision support query processing. And the other is data mining, specifically, association rule mining.As a high speed network, ATM technology has recently come to be a de facto standard. While other high performance network standards are also available, ATM networks are widely used from ...

13 Evaluation of remote backup algorithms for transaction-processing systems

Christos A. Polyzois, Héctor García-Molina

September 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 3

Publisher: ACM Press

Full text available: pdf(1.87 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A remote backup is a copy of a primary database maintained at a geographically separate location and is used to increase data availability. Remote backup systems are typically log-based and can be classified into 2-safe and 1-safe, depending on whether transactions commit at both sites simultaneously or first commit at the primary and are later propagated to the backup. We have built an experimental database system on which we evaluated the performance of the epoch and the dependency recons ...

Keywords: disaster recovery, hot spare, hot standby, remote backup

14 ReVive: cost-effective architectural support for rollback recovery in shared-memory multiprocessors

Milos Prvulovic, Zheng Zhang, Josep Torrellas
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2
Publisher: IEEE Computer Society, ACM Press
Full text available: pdf(1.38 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance, ReVive includes specialized hardware that performs frequent o ...

Keywords: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability

15 Link-based similarity: LSH forest: self-tuning indexes for similarity search
Mayank Bawa, Tyson Condie, Prasanna Ganesan
May 2005 **Proceedings of the 14th international conference on World Wide Web**
Publisher: ACM Press
Full text available: pdf(247.91 KB) Additional Information: full citation, abstract, references, index terms

We consider the problem of indexing high-dimensional data for answering (approximate) similarity-search queries. Similarity indexes prove to be important in a wide variety of settings: Web search engines desire fast, parallel, main-memory-based indexes for similarity search on text data; database systems desire disk-based similarity indexes for high-dimensional data, including text and images; peer-to-peer systems desire distributed similarity indexes with low communication cost. We propose an i ...

Keywords: peer-to-peer (P2P), similarity indexes

16 Concepts and paradigms of object-oriented programming
Peter Wegner
August 1990 **ACM SIGPLAN OOPS Messenger**, Volume 1 Issue 1
Publisher: ACM Press
Full text available: pdf(5.52 MB) Additional Information: full citation, abstract, citings, index terms

We address the following questions for object-oriented programming: *What is it? What are its goals? What are its origins? What are its paradigms? What are its design alternatives? What are its models of concurrency? What are its formal computational models? What comes after object-oriented programming?* Starting from software engineering goals, we examine the origins and paradigms of object-oriented programming, explore its language design alternativ ...

17 Epidemic algorithms in replicated databases (extended abstract)
D. Agrawal, A. El Abbadi, R. C. Steinke
May 1997 **Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**
Publisher: ACM Press
Full text available: pdf(1.59 MB) Additional Information: full citation, references, citings, index terms

18 Update propagation strategies to improve freshness in lazy master replicated databases

Esther Pacitti, Eric Simon

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(151.35 KB) Additional Information: full citation, abstract, index terms

Many distributed database applications need to replicate data to improve data availability and query response time. The two-phase commit protocol guarantees mutual consistency of replicated data but does not provide good performance. Lazy replication has been used as an alternative solution in several types of applications such as on-line financial transactions and telecommunication systems. In this case, mutual consistency is relaxed and the concept of freshness is used to measure the deviation ...

Keywords: Data replication, Distributed databases, Performance evaluation

19 Conflict detection tradeoffs for replicated data

Michael J. Carey, Miron Livny

December 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 4

Publisher: ACM Press

Full text available: pdf(2.50 MB) Additional Information: full citation, references, citings, index terms, review

Keywords: concurrency control, replicated data

20 Database replication with Slony-I

Ludovic Marcotte

June 2005 **Linux Journal**, Volume 2005 Issue 134

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(22.00 KB) Additional Information: full citation, abstract, index terms

Move up to a highly available cluster without leaving behind the open-source database you trust.

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next